

---

# Printed by EAST

---

UserID: bwong


Computer: WS10221

Date: 05/08/2008

Time: 15:40

## Document Listing

Document	Image pages	Text pages	Error pages
EP 1079570 A2	9	0	0
Total	9	0	0



Europäisches Patentamt  
European Patent Office  
Office européen des brevets



(11) **EP 1 079 570 A2**

(12) **EUROPEAN PATENT APPLICATION**

(43) Date of publication:  
**28.02.2001 Bulletin 2001/09**

(51) Int. Cl.<sup>7</sup>: **H04L 12/56, H04L 29/14**

(21) Application number: **00202894.2**

(22) Date of filing: **18.08.2000**

(84) Designated Contracting States:  
**AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU  
MC NL PT SE**  
Designated Extension States:  
**AL LT LV MK RO SI**

(30) Priority: **20.08.1999 US 378141**

(71) Applicant:  
**NORTEL NETWORKS CORPORATION**  
**Montreal Quebec H2Y 3Y4 (CA)**

(72) Inventors:  
• **Anderson, Loa**  
**125 33 Alvsjo (SE)**  
• **Felske, Kent**  
**Kantana, Ontario K2L 3V6 (CA)**  
• **Wang, Guo-Qiang**  
**Nepean, Ontario K2J 4L2 (CA)**

(74) Representative:  
**Ertl, Nicholas Justin**  
**Elkington and Fife,**  
**Prospect House,**  
**8 Pembroke Road**  
**Sevenoaks, Kent TN13 1XR (GB)**

(54) **Network data routing protection cycles for automatic protection switching**

(57) A computer network (22) processes data packets in the event of a network link failure. The network includes a plurality of routers (20) that deliver data packets to the network via a plurality of links (21). At least one router (20) includes a protection cycle manager (25). The protection cycle manager (25) has a protection cycle packet identifier (251) and a protection cycle packet processor (252). The protection cycle packet identifier (251) identifies, as protection cycle packets, data packets having a specific protection cycle format. The protection cycle packet processor (252) processes protection cycle packets to determine whether the packet destination corresponds to the routing node (20), and if the packet destination corresponds to the routing node (20), the protection cycle packet is treated by the routing node as a data packet received from the packet source via the failed link. Otherwise, if the packet destination does not correspond to the routing node (20), the protection cycle packet is sent to a protection cycle node for the routing node.




Figure 2

EP 1 079 570 A2

**Description**

**[0001]** The present invention relates to computer networks, and more specifically to a computer network that provides automatic protection switching to reroute data packets in the event of a network link failure.

**[0002]** In an Internet Protocol (IP) based computer network, data routing protocols such as Open Shortest Path First (OSPF), Intermediate System-Intermediate System (IS-IS), and Routing Information Protocol (RIP) are used to determine the path that data packets travel through the network. When a link between two network routers fails, the routing protocols are used to advertise the failure throughout the network. Most routers can detect a local link failure relatively quickly, but it takes the network as a whole a much longer time to converge. This convergence time is typically on the order of 10-60 seconds depending on the routing protocol and the size of the network. Eventually, all of the involved routers learn of the link failure and compute new routes for data packets to affected destinations. Once all the routers converge on a new set of routes, data packet forwarding proceeds normally. While the network is converging after a link fails, transient loops can occur which consume valuable bandwidth. Loop prevention algorithms have been proposed to eliminate such transient loops. When using these algorithms, routes are pinned until the network has converged and the new routes have been proven to be loop-free. Loop prevention algorithms have the advantage that data packets flowing on unaffected routes are not disrupted while transient loops are eliminated. The main drawback of loop prevention algorithms is that data packets directed out of a failed link get lost, or "black holed," during the convergence. Loop prevention algorithms also extend the convergence time somewhat while new routes are being verified to be loop-free.

**[0003]** According to a first aspect of the present invention, there is provided a protection cycle manager that processes data packets in the event of a link failure of a routing node that delivers data packets to a computer network via a plurality of links, the protection cycle manager comprising a protection cycle packet identifier that identifies, as protection cycle packets, data packets having a specific protection cycle format that includes a packet source and a packet destination, and a protection cycle packet processor that processes each protection cycle packet to determine whether the packet destination corresponds to the routing node wherein, if the packet destination corresponds to the routing node, the protection cycle packet is treated by the routing node as a data packet received from the packet source via the failed link, and wherein if the packet destination does not correspond to the routing node, the protection cycle packet is sent to a protection cycle node for the routing node.

**[0004]** According to another aspect of the present invention, there is provided a method of processing data

packets in the event of a link failure of a routing node that delivers data packets to a data network via a plurality of links, the method comprising the steps of:

identifying, as protection cycle packets, data packets having a specific protection cycle format that includes a packet source and a packet destination; and, processing each protection cycle packet to determine whether the packet destination corresponds to the routing node, wherein,

- (i) if the packet destination corresponds to the routing node, the protection cycle packet is treated by the routing node as a data packet received from the packet source via the failed link; and,
- (ii) if the packet destination does not correspond to the routing node, the protection cycle packet is sent to a protection cycle node for the routing node.

**[0005]** A preferred embodiment of the present invention includes an apparatus and method for processing data packets in the event of a link failure of a routing node that delivers data packets to a data network via a plurality of links. A protection cycle manager includes a protection cycle packet identifier and a protection cycle packet processor. The packet identifier identifies, as protection cycle packets, data packets having a specific protection cycle format that includes a packet source and a packet destination. The processor processes each protection cycle packet to determine whether the packet destination corresponds to the routing node, and if the packet destination corresponds to the routing node, the protection cycle packet is treated by the routing node as a data packet received from the packet source via the failed link. Otherwise, if the packet destination does not correspond to the routing node, the protection cycle packet is sent to a protection cycle node for the routing node.

**[0006]** A further embodiment may include a protection cycle packeter, that, in response to failure of a link for the routing node,, converts affected data packets routed over the failed link into protection cycle packets in the specific protection cycle format. The protection cycle manager may further advertise a link failure to the network using a routing protocol. The specific protection cycle format may include a label stack based on Multi-Protocol Label Switching (MPLS) which may include labels for the packet source and the packet destination. The protection cycle node may be a next node on a Label Switched Path (LSP), in which case, the LSP may be based on network topology information which may be derived from a network protocol.

**[0007]** A preferred embodiment also includes a data router that delivers data packets to a data network via a plurality of links, the router processing data pack-

ets in the event of a link failure. The router includes a data interface for data packets to enter and exit the router, and a protection cycle manager. The protection cycle manager includes a protection cycle packet identifier and a protection cycle processor. The packet identifier identifies, as protection cycle packets, data packets having a specific protection cycle format that includes a packet source and a packet destination. The processor processes protection cycle packets to determine whether the packet destination corresponds to the routing node, and: (i) if the packet destination corresponds to the router, the protection cycle packet is treated by the router as a data packet received from the packet source via the failed link, and, (ii) if the packet destination does not correspond to the routing node, the protection cycle packets are sent to a protection cycle node for the router.

**[0008]** Another preferred embodiment includes a computer network having a plurality of data packet streams. The network includes a plurality of subnetworks, each subnetwork having at least one application that generates a stream of data packets for transmission over the computer network, and a plurality of routers that deliver data packets to the network via a plurality of links. At least one router processes data packets in the event of a link failure. The at least one router includes a plurality of data interfaces for streams of data packets to enter and exit the at least one router, and a protection cycle manager having a protection cycle packet identifier and a protection cycle processor. The packet identifier identifies, as protection cycle packets, data packets having a specific protection cycle format that includes a packet source and a packet destination. The processor processes protection cycle packets to determine whether the packet destination corresponds to the routing node, and: (i) if the packet destination corresponds to the at least one router, the protection cycle packet is treated by the at least one router as a data packet received from the packet source via the failed link, and, (ii) if the packet destination does not correspond to the routing node, the protection cycle packets are sent to a protection cycle node for the at least one router.

**[0009]** A preferred embodiment includes a computer program product for use on a computer system for processing data packets in the event of a link failure of a routing node that delivers data packets to a data network via a plurality of links. The computer program product comprises a computer-usable medium having computer-readable program code thereon. The computer readable program code includes program code for identifying, as protection cycle packets, data packets having a specific protection cycle format that includes a packet source and a packet destination. There is also program code for processing protection cycle packets to determine whether the packet destination corresponds to the routing node, and: (i) if the packet destination corresponds to the routing node, the protection cycle

packet is treated by the routing node as a data packet received from the packet source via the failed link, and, (ii) if the packet destination does not correspond to the routing node, the protection cycle packets are sent to a protection cycle node for the routing node.

**[0010]** The present invention will be more readily understood by reference to the following detailed description taken with the accompanying drawings, in which:

Figure 1 is an illustration of a computer network which uses MPLS p-cycle automatic protection switching according to a preferred embodiment of the present invention;

Figure 2 is an illustration of a network node router which supports p-cycles according to a another preferred embodiment of the present invention; and,

Figure 3 is a flow chart illustration of the logical steps in a method of providing automatic protection switching according to a further preferred embodiment of the present invention.

**[0011]** A preferred embodiment of the present invention uses Multi-Protocol Label Switching (MPLS) with explicit routing to establish an MPLS layer protection cycle (p-cycle) that provides automatic protection switching to reroute data packets in the event of a network link failure. An MPLS Label Switched Path (LSP) tunnel passes through the end points of the link that will be protected. The LSP tunnel established in this way forms an MPLS p-cycle through which failed-link packets are directed. A given p-cycle may protect one or multiple links. In a preferred embodiment, a standard multi-layered MPLS label stack is used in the p-cycle. In other embodiments, a p-cycle may be implemented in another protocol other than MPLS.

**[0012]** The p-cycle may be configured by hand, or automatically established using network link state and topology information derived from a routing protocol, such as Open Shortest Path First (OSPF). Various algorithms may be used to automatically compute the specific structure of a given p-cycle. A preferred embodiment uses a network management application for this purpose since the p-cycle will not follow an optimal path. When a p-cycle is bi-directional, a bandwidth protection mechanism may be implemented so that some of the p-cycle traffic goes one way, and the rest the other way.

**[0013]** A preferred embodiment uses a diffusion-based loop prevention algorithm as is known in the art to detect when the network has converged so that it is safe to switch to the new routed path. Such diffusion algorithms are discussed, for example, in Garcia-Lunes-Aceves, J.J., "Loop-Free Routing Using Diffusing Computations," *IEEE/ACM Transactions on Networking*, vol. 1, no. 1, 1993, pp. 130-141, which is hereby incorporated herein by reference. Using p-cycles with a loop

prevention algorithm allows for uninterrupted service in the event of a link failure without black holing of packets on the failed link.

[0014] Figure 1 is an illustration of a computer network using a preferred embodiment of the present invention for MPLS p-cycle automatic protection switching. In Figure 1, network nodes A 10 through H 17 are normally linked via the dashed lines to route data packets from node to node. In the network segment shown, the normal network links between A and E, A and F, B and D, B and E, and H and F are also protected by a p-cycle 18. The p-cycle 18 forms a loop through the network so that a packet sent through the p-cycle will eventually come back to its origination node if not taken out of the p-cycle by one of the nodes it passes through. In Figure 1, LSP p-cycle 18 traverses the network segment from node A 10 to B 11 to C 12 ... to G 16 to H 17 and back to A 10.

[0015] Only one p-cycle 18 is shown in Figure 1. In practice, a p-cycle is established for every set of links to be protected. A preferred embodiment can operate successfully in any arbitrary network topology. It should be noted, however, that to realize full link-level protection, for every two neighbors A and B connected by link L in the network, another network path between A and B must exist that does not include L.

[0016] Various options may be employed with respect to network-level encapsulation on the original link. For example, the original network-layer encapsulation (e.g., IP) may be tunneled in the p-cycle LSP. If MPLS is being used on the original link, then the labeled packet may be tunneled on the backup link using MPLS label stacking. Multiple independent link failures may be tolerated using multiple layers of tunneling.

[0017] The router for each node monitors its own local links. When a link failure is detected, the router for an affected node quickly routes the data packet traffic over to the p-cycle 18. Then, the network routing protocol advertises the link failure so that the network can be re-routed without the failed link, and a loop-prevention mechanism determines that the re-routed network is loop-free. Packet traffic may then be switched to the re-routed network and new p-cycles recalculated as necessary.

[0018] Figure 2 is an illustration of a network node router which supports p-cycles according to a preferred embodiment. Network node router 20 is a part of a computer network 22 of routers in mutual communication via a plurality of network node data links 21. Router 20 also serves to connect one or more local area networks (LANs) 23 having one or more workstations 231. Data packets enter and exit the router 20 as controlled by a data interface driver 24 which is connected to the network node links 21.

[0019] Figure 3 is a flow chart illustration of the logical steps in a corresponding method of providing automatic protection switching according to a preferred embodiment. A protection cycle manager 25 includes a

p-cycle packet identifier 251 that, in step 301, identifies as p-cycle packets, data packets that have a p-cycle label stack, which, in a preferred embodiment, is a standard MPLS label stack. In such an embodiment, the top label in the stack indicates the next node in the p-cycle. The next label on the stack is the identity of the destination node that ultimately receives the packet, the third label in the stack is the identity of the node creating the label stack. Identified p-cycle packets are processed by p-cycle packet processor 252 which, in step 302, pops the topmost label off the label stack and checks the next label to see if the router node's own identity is in the destination node position in the label stack. If not, the label for the next p-cycle node is pushed onto the stack and the packet is sent by the data interface driver 24 via the node link 21 to the next node on the p-cycle, step 303. If, in step 302, the router node's own identity is carried in the destination node position in the label stack, the source node label in the label stack is checked to determine which network link the packet normally would have used, step 304. The p-cycle label stack is then deleted, step 305, and thereafter, the packet is treated as if it had been received via the normal network link 21 from the source node, step 306.

[0020] In a preferred embodiment, the network node router also includes a network link monitor 26 in communication with the data interface driver 24. When the link monitor 26 detects a failed link, step 307, protection cycle packeter 253 attaches to affected data packets a p-cycle label stack having appropriate labels for source node, destination node, and p-cycle node, step 308, and the p-cycle packets are then sent to the p-cycle node for that router, step 303. A link failure also is advertised to the network using the routing protocol, step 309. A new network route is then established to replace the failed link, step 310, and a loop prevention algorithm is used to determine that the new network routes have converged and are loop-free, step 311.

[0021] Preferred embodiments of the invention, or portions thereof (e.g., the p-cycle packet processor 252, p-cycle packet identifier 251, link monitor 26, etc.), may be implemented in any conventional computer programming language. For example, preferred embodiments may be implemented in a procedural programming language (e.g., "C") or an object oriented programming language (e.g., "C++" or "JAVA"). Alternative embodiments of the invention may be implemented as preprogrammed hardware elements (e.g., application specific integrated circuits), or other related components.

[0022] Alternative embodiments of the invention may be implemented as a computer program product for use with a computer system. Such implementation may include a series of computer instructions fixed either on a tangible medium, such as a computer readable media (e.g., a diskette, CD-ROM, ROM, or fixed disk), or transmittable to a computer system via a modem or other interface device, such as a communica-

tions adapter connected to a network over a medium. The medium may be either a tangible medium (e.g., optical or analog communications lines) or a medium implemented with wireless techniques (e.g., microwave, infrared or other transmission techniques). The series of computer instructions preferably embodies all or part of the functionality previously described herein with respect to the system. Those skilled in the art should appreciate that such computer instructions can be written in a number of programming languages for use with many computer architectures or operating systems. Furthermore, such instructions may be stored in any memory device, such as semiconductor, magnetic, optical or other memory devices, and may be transmitted using any communications technology, such as optical, infrared, microwave, or other transmission technologies. It is expected that such a computer program product may be distributed as a removable medium with accompanying printed or electronic documentation (e.g., shrink wrapped software), preloaded with a computer system (e.g., on system ROM or fixed disk), or distributed from a server or electronic bulletin board over the network (e.g., the Internet or World Wide Web).

**[0023]** Although various exemplary embodiments of the invention have been disclosed, it should be apparent to those skilled in the art that various changes and modifications can be made that will achieve some of the advantages of the invention without departing from the true scope of the invention. These and other obvious modifications are intended to be covered by the appended claims.

#### Claims

1. A protection cycle manager (25) that processes data packets in the event of a link failure of a routing node (20) that delivers data packets to a computer network (22) via a plurality of links (21), the protection cycle manager (25) comprising a protection cycle packet identifier (251) that identifies, as protection cycle packets, data packets having a specific protection cycle format that includes a packet source and a packet destination, and a protection cycle packet processor (252) that processes each protection cycle packet to determine whether the packet destination corresponds to the routing node (20) wherein, if the packet destination corresponds to the routing node (20), the protection cycle packet is treated by the routing node (20) as a data packet received from the packet source via the failed link, and wherein if the packet destination does not correspond to the routing node (20), the protection cycle packet is sent to a protection cycle node for the routing node (20).
2. A protection cycle manager (25) according to claim 1, further comprising a protection cycle packeter (253), that converts, in response to failure of a link for the routing node (20), affected data packets routed over the failed link into protection cycle packets in the specific protection cycle format.
3. A protection cycle manager (25) according to claim 1 or 2, wherein the protection cycle manager (25) is arranged to advertise a link failure to the network (22) using a routing protocol.
4. A protection cycle manager (25) according to claim 1, 2 or 3, wherein the specific protection cycle format includes a label stack based on Multi-Protocol Label Switching (MPLS).
5. A protection cycle manager (25) according to claim 4, wherein the label stack includes labels for the packet source and the packet destination.
6. A protection cycle manager (25) according to any preceding claim, wherein the protection cycle node is a next node on a Label Switched Path (LSP).
7. A protection cycle manager (25) according to claim 6, wherein the LSP is based on network topology information.
8. A protection cycle manager (25) according to claim 7, wherein the network topology information is derived from a network protocol.
9. A protection cycle manager (25) according to any preceding claim, further comprising means for establishing a new network route to replace a failed link.
10. A protection cycle manager (25) according to any preceding claim, further comprising means arranged to apply a loop prevention algorithm after a link failure for determining that the network (22) has converged and is loop-free.
11. A data router (20) arranged to deliver data packets to a computer network (22) via a plurality of links (21), the router processing data packets in the event of a link failure, the router (20) comprising a data interface (24) for data packets to enter and exit the router (20) and a protection cycle manager (25) according to any preceding claim, wherein the router is a routing node.
12. A computer network (22) having a plurality of data packet streams, the network comprising a plurality of subnetworks (23), each subnetwork having at least one application that generates a stream of data packets for transmission over the computer network (22) and a plurality of routers (20) that deliver data packets to the network via a plurality of links, at least one router processing data packets in

the event of a link failure, the at least one router including a plurality of data interfaces for streams of data packets to enter and exit the at least one router and a protection cycle manager (25) according to any of claims 1 to 8, the at least one router comprising a routing node.

13. A method of processing data packets in the event of a link failure of a routing node that delivers data packets to a data network via a plurality of links, the method comprising the steps of:

identifying, as protection cycle packets, data packets having a specific protection cycle format that includes a packet source and a packet destination; and,  
processing each protection cycle packet to determine whether the packet destination corresponds to the routing node, wherein,

- (i) if the packet destination corresponds to the routing node, the protection cycle packet is treated by the routing node as a data packet received from the packet source via the failed link; and,  
(ii) if the packet destination does not correspond to the routing node, the protection cycle packet is sent to a protection cycle node for the routing node.

14. A method according to claim 13, further comprising the step of converting, in response failure of a link for the routing node, affected data packets routed over the failed link into protection cycle packets in the specific protection cycle format.

15. A method according to claim 13 or 14, further comprising the step of advertising a link failure to the network using a routing protocol.

16. A method according to claim 13, 14 or 15, further comprising the step of establishing a new network route to replace a failed link.

17. A method according to any of claims 13 to 16, further comprising the step of using a loop prevention algorithm after a link failure for determining that the network has converged and is loop-free.

18. A method according to any of claims 13 to 17, wherein the specific protection cycle format includes a label stack based on Multi-Protocol Label Switching (MPLS).

19. A method according to claim 18, further comprising the step of including labels for the packet source and the packet destination within the label stack.

20. A method according to any of claims 13 to 19, wherein the protection cycle node is a next node on a Label Switched Path (LSP).

21. A method according to claim 20, wherein the LSP is based on network topology information.

22. A method according to claim 21, further comprising the step of deriving the network topology information from a network protocol.

23. A computer program comprising computer program code means adapted to execute the protection cycle manager of any of claims 1 to 10 when the program is run on a computer.

24. A computer program comprising computer program code means adapted to perform all the steps of any one of claims 13 to 22 when the program is run on a computer.




Figure 1






Figure 2

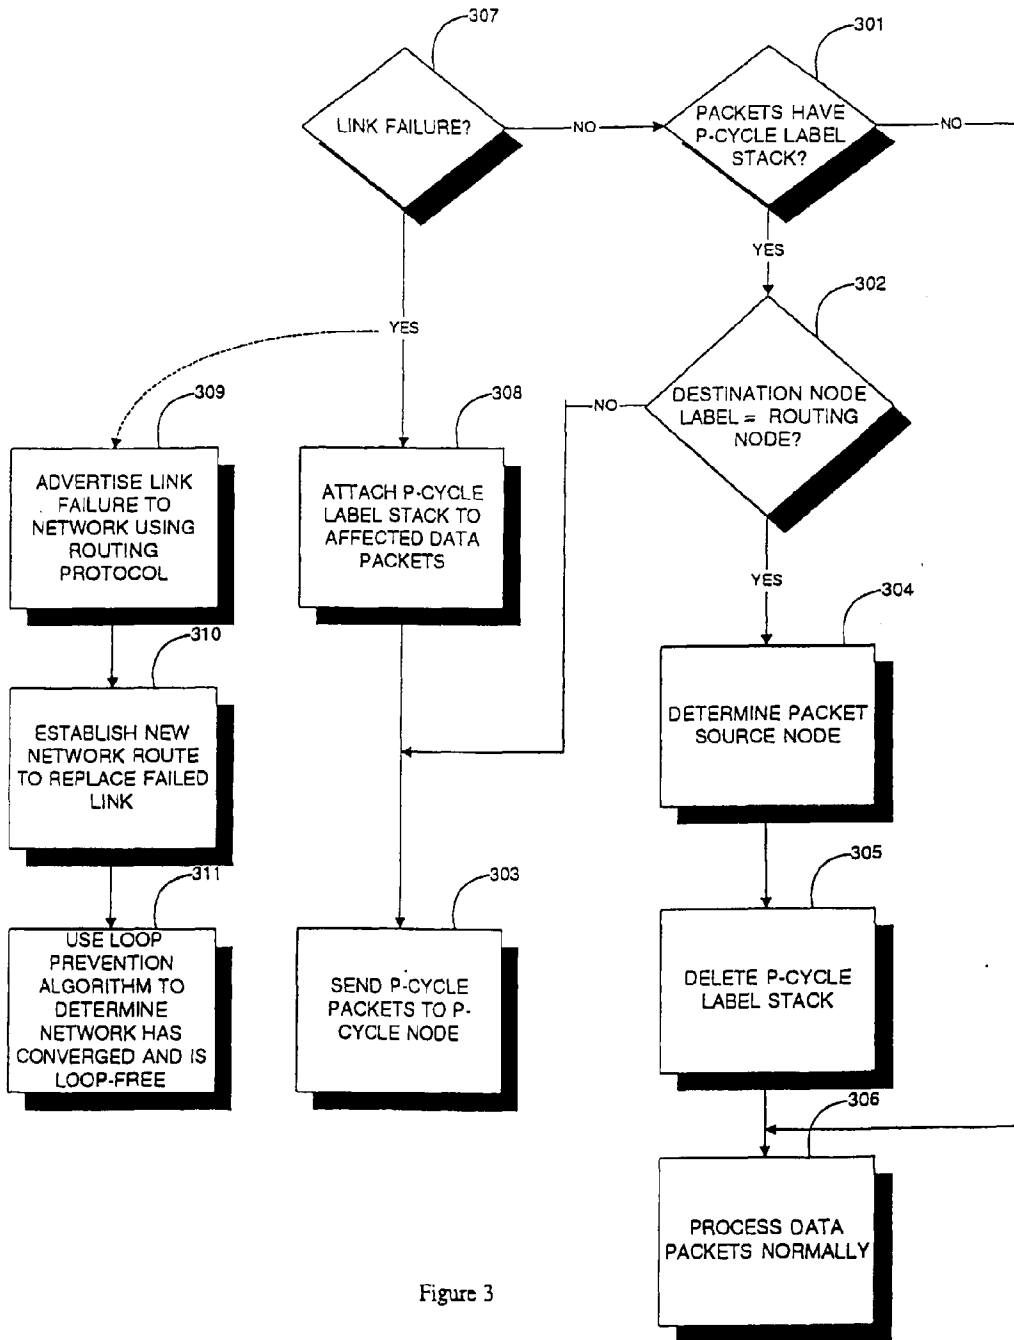


Figure 3